

WE CLAIM:

1. A method in a computer system for aggregating storage in a data storage network having a server with one or more consumers, a storage system with available storage,  
5 and a communication fabric linking the server and the storage system, comprising:

pooling the available storage to create virtual drives;  
presenting the virtual drives to the server over the fabric;

10 in response, receiving a logical command from the server for access to the available storage represented by the virtual drives; and

transmitting the logical command to a controller of the available storage identified in the logical command.

15 2. The method of Claim 1, further including establishing a direct communication link over the communication fabric between the server and the storage controller to perform actions indicated in the logical command.

20 3. The method of Claim 1, wherein the virtual drives comprise a combination of logical unit numbers (LUNs).

4. The method of Claim 3, wherein the pooling includes dividing volumes in the available storage into pages and wherein the combination comprises one or more  
25 volumes of LUN pages based on the available storage pages.

5. The method of Claim 4, further including determining if logical command is a write to one of the pages marked as snapped and based on such determination,

allocating a new page for the marked page and clearing the snapped marking.

6. The method of Claim 5, wherein the logical command is a small computer system interface (SCSI) read or write.

5 7. The method of Claim 1, further including receiving a snapshot configuration command and copying the virtual drives included in the snapshot configuration command to create snapshot virtual drives.

10 8. A data storage network with virtualized data storage, comprising:

a communication fabric;

a server system linked to the communication fabric and running applications that transmit data access commands over the communication fabric;

15 a storage system linked to the communication fabric including data storage devices and a controller for managing access to the data storage devices; and

a storage aggregator linked to the communication fabric having virtual drives comprising a logical representation of the data storage devices, wherein the storage aggregator receives the data access commands pertaining to the virtual drives and forwards the data access commands to the controller of the storage system.

25 9. The data storage network of Claim 8, wherein the storage system responds to the forwarded data access commands by communicating with the server system over the fabric.

10. The data storage network of Claim 8, wherein the storage aggregator includes a virtual controller adapted for at least periodically creating the virtual drives based on available volumes in the data storage devices.

5 11. The data storage network of Claim 10, wherein the creating includes dividing the available volumes into pages and wherein the virtual drives include aggregate volumes comprising a plurality of logical unit number (LUN) pages representing the pages from the available volumes.

10 12. The data storage network of Claim 8, wherein the data access commands are SCSI commands.

13. The data storage network of Claim 8, wherein data access commands are SCSI RDMA Protocol (SRP) commands.

15 14. The data storage network of Claim 8, wherein the communication fabric is a switched matrix having at least one switch and the server system, the storage system, and the storage aggregators each include an adapter linked to the switch for transmitting and receiving messages to and from the switch.

20 15. The data storage network of Claim 14, wherein the server system adapter is a host channel adapter and the storage system and storage aggregator adapters are target channel adapters.

25 16. The data storage network of Claim 15, wherein the communication fabric is an InfiniBand Architecture switched communication fabric.

17. A storage aggregation apparatus for virtualizing data access commands, comprising:

an input and output interface linking the storage aggregation apparatus to a digital data communication fabric, wherein the interface receives a data access command from a host server over the fabric;

a command processor configured to parse a data movement portion from the data access command and to transmit the data movement portion to a storage controller; and

a mechanism for receiving a reply signal from the storage controller in response to acting on the received data movement portion directly with the host server and for transmitting a data access response to the host server via the interface and fabric based on the storage controller reply signal.

18. The apparatus of Claim 17, wherein the data access command is a SCSI command and the data access response is a SCSI response.

19. The apparatus of Claim 17, wherein the apparatus is adapted for controlling access to a pool of storage and further includes a mechanism for creating virtual storage based on the storage pool and for presenting the virtual storage to the host server over the fabric.

20. The apparatus of Claim 17, wherein the apparatus is a node on the fabric remote from the storage controller.

21. The apparatus of Claim 17, wherein the apparatus is a component of a storage system including the storage controller.